

Non-Gaussian Posteriors arising from Marginal Detections

Bruce A. Bassett

South African Astronomical Observatory, Observatory, Cape Town, South Africa

*Department of Mathematics and Applied Mathematics, University of Cape Town,
Rondebosch, 7700, Cape Town, South Africa*

African Institute for Mathematical Sciences, Muizenberg, South Africa

*Centre for High Performance Computing, 15 Lower Hope St, Rosebank, Cape Town, South
Africa*

*Perimeter Institute for Theoretical Physics, 31 Caroline St. N., Waterloo, ON, N2L
2Y5, Canada*

Niayesh Afshordi

*Perimeter Institute for Theoretical Physics, 31 Caroline St. N., Waterloo, ON, N2L
2Y5, Canada*

*Department of Physics and Astronomy, University of Waterloo, 200 University Avenue
West, Waterloo, ON, N2L 3G1, Canada*

ABSTRACT

We show that in cases of marginal detections ($\sim 3\sigma$), such as that of Baryonic Acoustic Oscillations (BAO) in cosmology, the often-used Gaussian approximation to the full likelihood is very poor, especially beyond ($\sim 3\sigma$). This can radically alter confidence intervals on parameters and implies that one cannot naively extrapolate 1σ error bars to 3σ and beyond. We propose a simple fitting formula which corrects for this effect in posterior probabilities arising from marginal detections. Alternatively the full likelihood should be used for parameter estimation rather than the Gaussian approximation of a just mean and an error.

Subject headings: Methods: statistical, Cosmology, cosmological parameters, large-scale structure of Universe

Making observational detections in cosmology is hard. Two to three sigma evidence for new physics is common in the literature and the community is rightly skeptical of results with marginal statistical significance. Even in the case of 5σ results, detections can be questioned if the results are surprising or conflicting data is released.

When detections are expected there is a natural tendency to accept them more readily. Baryon Acoustic Oscillations (BAO) (BAO; *e.g.*, Bassett & Hlozek 2010) are a good example. The original BAO results of the SDSS and 2dF teams (Eisenstein et al. 2005; Cole et al. 2005) were at less than 3σ significance. Nevertheless, the detection has been essentially unanimously accepted by the community (although see Sylos Labini et al. 2009) despite the difficulty of localizing the BAO peak, *e.g.*, illustrated by the shift in BAO results between 2007 and 2009 in Percival et al. (2007) and Percival et al. (2010) and recent studies of mock catalogs (Kazin et al. 2010b) which suggest that the BAO peak would be invisible in at least 10% of SDSS DR7-sized samples.

Such detections are easy to accept since the detected peak is precisely in the place where it was expected to be. This willingness to accept marginal detections carries two dangers. First, the signal may actually be pure statistical fluctuation and hence provide precise but inaccurate knowledge. Second, it can actively discourage publication of other studies which are apparently at odds with the ‘detection’.

In this *letter*, we point out that the possibility that the detection is not real, but just a statistical fluctuation due to noise, has a significant impact on use of the data for parameter estimation. This arises because the full posterior typically has very non-gaussian tails even if the likelihood is gaussian, as illustrated in Figure (1). Like a scorpion, the sting is in the tail.

The prototypical example we have in mind is the use of the BAO peak for cosmology, but the principle is valid generally. While the current best BAO combined detection is at the $\sim 4.9\sigma$ level from the SDSS, 6df and WiggleZ catalogues (Blake et al. 2011), lower significance detections will always mark the frontier of the subject as we push to higher redshifts and different samples. The search for BAO with the SDSS, 2QZ and 2SLAQ survey quasars illustrates our point. While they show a peak in the expected place ($105h^{-1}\text{Mpc}$, they also show a (presumably) fake peak near $85h^{-1}\text{Mpc}$ (Sawangwit et al. 2011). A more standard example is the recent 6df survey which found 2.4σ evidence for BAO while the WiggleZ detection is 3.2σ at $z = 0.6$. This will continue with a number of BAO first detections still to come in the next few years:

- The first separate detections of the radial and transverse BAO peaks which will yield $H(z)$ and $d_A(z)$ separately. There is a claim of detection of the radial BAO (Gaztañaga et al.

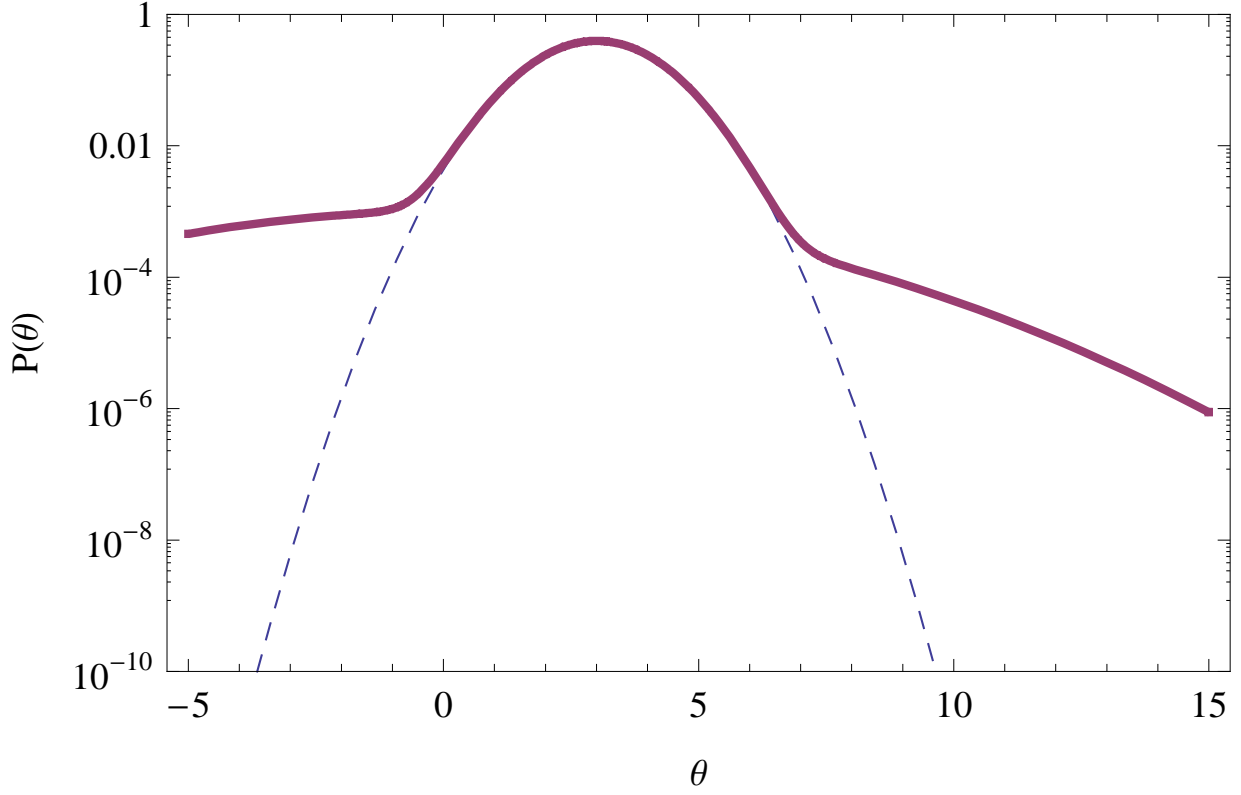


Fig. 1.— An example of how the posterior probability distribution has extended non-Gaussian wings, for finite detection probability: A peak is detected at 99% confidence level ($P_{detect} = 0.99$) at $\theta = 3 \pm 1$, while the gaussian prior of $\theta = 0 \pm 4$ was assumed. The dashed curve shows the naive gaussian posterior for θ , which *e.g.*, would suggest $\theta = 7$ is ruled out at 4σ level, while it is actually allowed at 99.5% confidence level.

2009). However this is somewhat controversial and the associated uncertainty illustrates the main points of this paper (Miralda-Escude 2009; Kazin et al. 2010a).

- The first detection of BAO in photometric redshift surveys.¹
- The first detection of BAO in cluster data. The current status is the 2-2.5 σ evidence from the maxBCG cluster catalogue (Hütsi 2010).
- The first detection of BAO in neutral hydrogen, HI.
- The first detection of BAO in the Lyman- α forest (McDonald & Eisenstein 2007). This is a method that will be employed by both the BOSS and LAMOST surveys.
- The first $z > 2$ BAO detection with Lyman break galaxies.

In addition, as future surveys progress, it will be tempting to split a given volume up into narrower redshift bins to provide more data points. In doing so it is standard to ignore the statistical significance of the BAO detections in forecasting the resulting constraints on dark energy, just as it has been standard to ignore them in using current BAO results: prominent examples include the WMAP7 analysis (Komatsu et al. 2011), the SDSS supernova survey (Kessler et al. 2009) and the grid marginalisation component of the Supernova Legacy (SNLS) 3-year analysis (Sullivan et al. 2011), all of which use the gaussian approximation to SDSS BAO results.

As a toy model to illustrate the impact of ignoring the finite detection probability of the peak, consider the posterior probability derived from a BAO experiment, $P(\theta|d, detected)$ where d is the data, θ are the parameters to be estimated (*e.g.*, w_0, w_a, Ω_k) and *detected* signifies the assumption that the apparent peak is not just noise, *i.e.*, the underlying model is correct. Then the full posterior is (Press 1997; Kunz et al. 2007):

$$P(\theta|d) = P_{detect}P(\theta|d, detected) + (1 - P_{detect})P(\theta), \quad (1)$$

where P_{detect} is the statistical significance of the detection and $P(\theta)$ is the prior probability for the parameters which coincides with $P(\theta|d, noise)$, the knowledge gained if the apparent peak is assumed to be pure noise, since in this case there is no new information. One of the most important implications of this formula is that the resulting posterior is very non-gaussian, with catastrophic widening where the likelihood drops below the prior. An example

¹Current surveys such as Mega-Z lack the number density to reduce shot noise to a level where detection is possible. DES (The Dark Energy Survey Collaboration 2005) and PanSTARRS (Cai et al. 2009) should provide the first detections while LSST will provide exquisite results (Zhan et al. 2009).

of this non-gaussian distribution is shown in Figure 1. An important component of the full posterior is that it must be correctly normalized: $\int P(\theta|d)d\theta = 1$.

Assuming that the prior is much wider than the measurement, *i.e.*, $P(\theta|d, detected) \gg P(\theta)$ near the peak of the likelihood, Figure (2) shows how the errorbars grow as one decreases the detection probability P_{detect} . This is done by finding the 68.3%, 95.5%, and 99.7% regions, using the normalization $P_{detect} < 1$ for a gaussian distribution.

Let us look at the example of $P_{detect} = 0.99$ (Figure 1). Although the 1 and 2- σ error bars are essentially unaffected, these are of little interest in parameter estimation since any real conclusions must be supported at the 99.7% (3σ) level or better. At this level, Figure (2) shows that there is a dramatic change: there are *no* constraints on the parameter at all at this significance, despite the likelihood $P(\theta|d, detected)$ possibly claiming wonderful constraints. What is happening is that one is transitioning from the likelihood to the prior when one moves sufficiently far from the maximum likelihood value such that:

$$P(\theta|d, detected) = \frac{(1 - P_{detect})}{P_{detect}} P(\theta) \simeq (1 - P_{detect}) P(\theta). \quad (2)$$

Since the prior should vary with θ much less rapidly than the likelihood, there is a point at which one's constraints are actually driven by the prior, not the data. The basic concept is simple: one cannot extrapolate data beyond its realm of validity, as illustrated by the figures.

We can go further and show how the non-gaussian sting can radically change even 1- σ errors when combining constraints from different measurements. In Fig. (3) we show a toy model with two independent measurements of two parameters $(D_1, D_2) = (0.3, 0.3) \pm (0.1, 0.15)$ and $(D_1, D_2) = (0.7, 0.7) \pm (0.1, 0.15)$, which are different due to some unaccounted for systematic error. Furthermore, each measurement is based on a 3.6σ detection, which causes a slight expansion of 1 and 2σ constraints (blue solid contours, corresponding to $\Delta\chi^2 = 2.3$ and 6) with respect to the gaussian constraints (black dotted contours), as described in this *letter* (see Equation 3 below). However, as we see on the second panel, combining the two constraints leads to radically different posteriors, with and without the assumption of gaussianity, *i.e.*, even the non-gaussian 1σ combined errors are much bigger than the gaussian ones, illustrating that the effects need not only limited be limited to $> 3\sigma$ results.

The frequentist analog of this statement is that, for a finite detection confidence level, the function $\chi^2(\theta|d)$ should asymptote to a plateau far away from its minimum, rather than growing indefinitely. The maximum difference between the plateau and the minimum is roughly the square of the signal-to-noise of the detection, S/N . Therefore, the gaussian (quadratic) approximation to the likelihood (χ^2) breaks down when $\Delta\chi_{\text{gauss}}^2 \gtrsim (S/N)^2$.

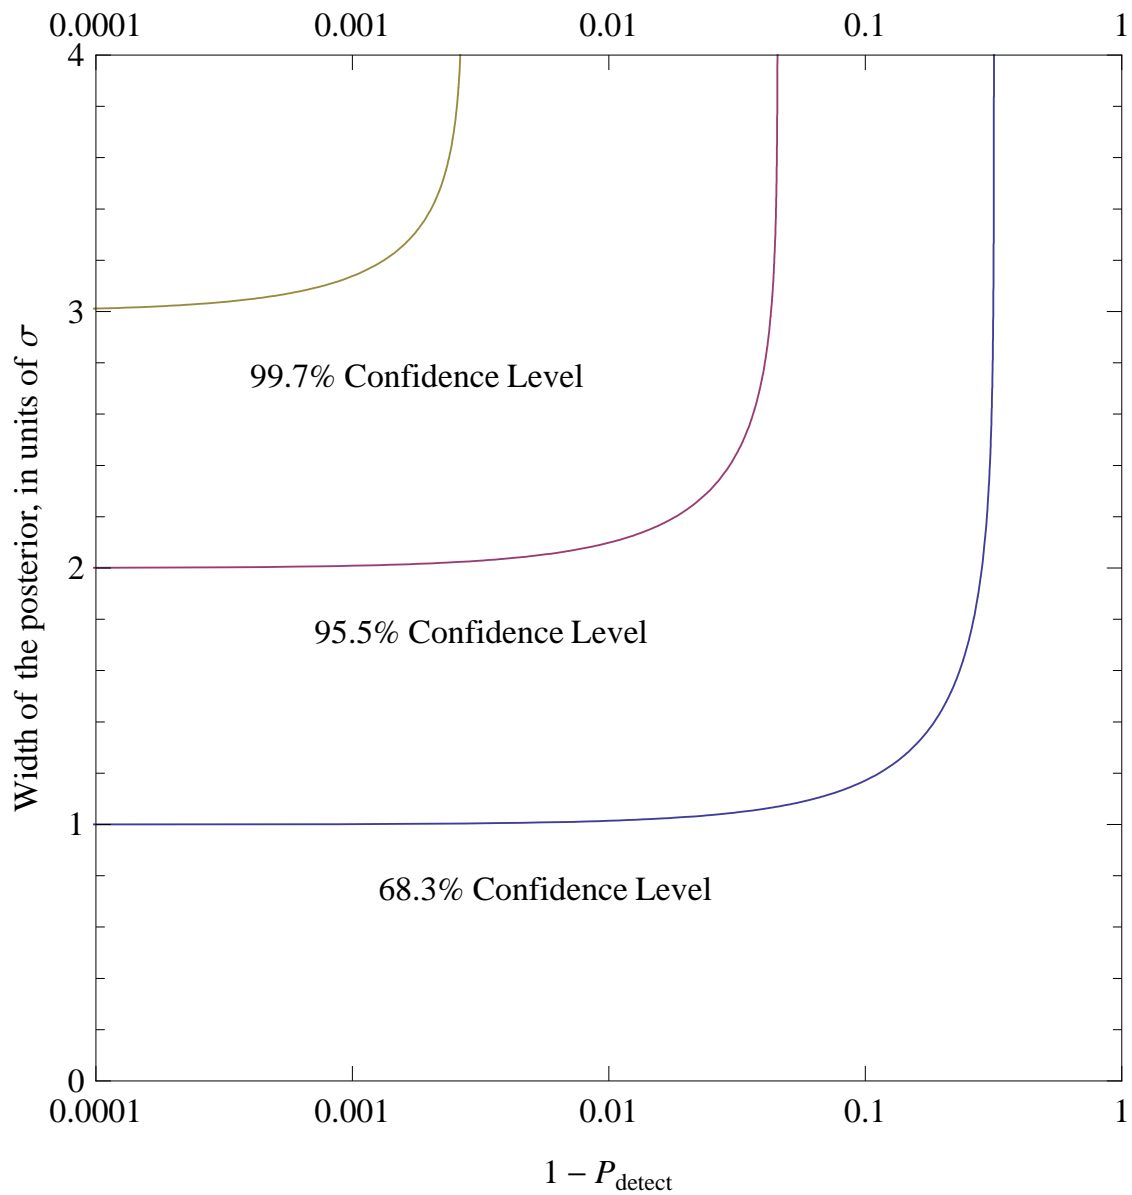


Fig. 2.— This figure shows how 68.3%, 95.5%, and 99.7% confidence levels for the width of the posterior probability distribution (associated with 1, 2, and 3σ gaussian errors) rapidly expand with decreasing detection probability, P_{detect} .

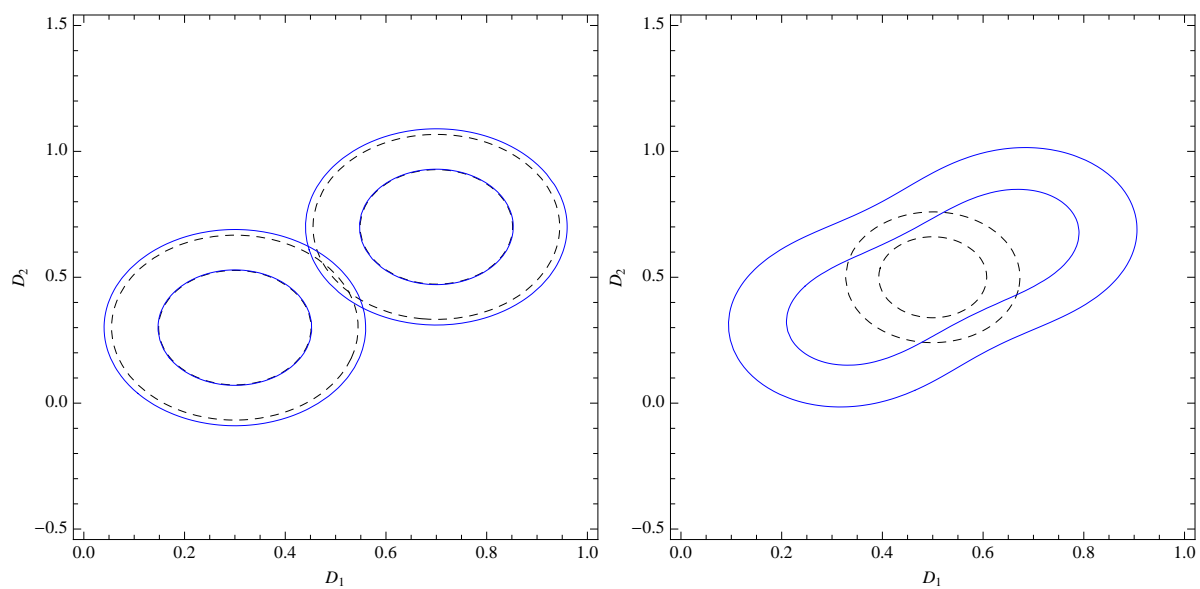


Fig. 3.— A toy example of how finite detection significance can affect even 1σ combined constraints: On the left, the solid (dashed) contours show the non-gaussian (gaussian) 1 and 2σ constraints based on two measurements of D_1 and D_2 , which each have 3.6σ significance. The right panel shows the combined constraints, which are radically different with and without the assumption of gaussianity.

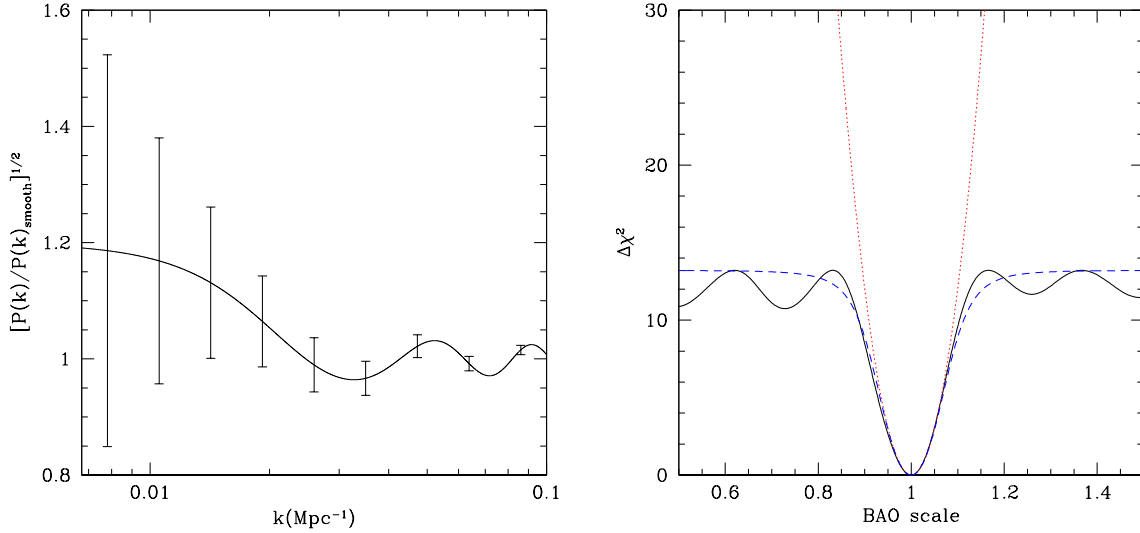


Fig. 4.— *Left:* (a) Power spectrum of a hypothetical galaxy survey, normalized to a smooth power spectrum without BAO oscillations (see text for details). *Right:* (b) $\Delta\chi^2$ ($= -2 \log[\text{likelihood}]$) for the position of the BAO peak in the linear power spectrum ($k < 0.1 \text{ Mpc}^{-1}$) for this survey. The red dotted curve shows the quadratic/gaussian approximation which grows indefinitely, while the actual $\Delta\chi^2$ (black solid curve) saturates at $\sim (S/N)^2 = 13.1$ for the detection. The blue dashed curve is an interpolation between the gaussian and the asymptotic regime that we suggest in Eq. (3) and is a good fit to the actual $\Delta\chi^2$.

As an example, in Figure (4) we consider the expected constraints on the BAO scale from a hypothetical galaxy survey with 10^6 galaxies and bias of 1.5, spread over a volume of 0.8 Gpc^3 at small redshifts ². In Figure (4a), we show the expected power spectrum+errors, normalized to a smooth power spectrum without BAO oscillations, using the Eisenstein & Hu (1999) fit to the transfer function. The solid curve in Figure (4b) shows $\Delta\chi^2$ for fitting this spectrum with a different BAO (or sound horizon) scale, marginalizing over the amplitude of the oscillations. Here, we only include linear scales, conservatively defined as $k < 0.1 \text{ Mpc}^{-1}$, and assume gaussian errors for $\Delta(k) = [k^3 P(k)/(2\pi^2)]^{1/2}$.

While the gaussian approximation to the likelihood (dashed curve in Figure 4b), is a good approximation close to the minimum, it grows indefinitely, while the real $\Delta\chi^2$ saturates at $\sim (S/N)^2$. In order to reflect this, we propose a simple analytic function to approximate the true difference between χ^2 and its minimum value:

$$\Delta\chi^2 \approx \frac{\Delta\chi_{\text{gauss}}^2}{[1 + (S/N)^{-4}\Delta\chi_{\text{gauss}}^4]^{1/2}}. \quad (3)$$

As shown in Figure (4b), this takes the quadratic shape of the gaussian approximation close to the minimum since the denominator is then negligible, but Eq. (3) guarantees that $\Delta\chi^2$ remains smaller than $(S/N)^2$ of the detection, far from its minimum, which limits the statistical power of low signal-to-noise detections in constraining parameters. While the interpolating function (3) is in good agreement with the actual $\Delta\chi^2$, we should note that it is only an approximation, and ideally one should use the full likelihood of the model fitting the data (*e.g.*, the full galaxy power spectrum) for an accurate statistical analysis.

Finally, Figure (5) demonstrates the effect of non-gaussian posteriors on cosmological constraints. Here, we compare the gaussian approximation to likelihood distribution for distances to $z = 0.2$ and $z = 0.35$ in Sloan Digital Sky Survey (SDSS; Percival et al. 2010), with our expectation from Equation (3). Given that total S/N for BAO detection in Percival et al. (2010) is $\sqrt{13.1}$ for two degrees of freedom, we see significant deviations from gaussian likelihoods beyond 99% confidence level ³.

Conclusions – In modern cosmology, as in many other branches of science, there is a clear division of labour between theorists and experimentalists. The theorists often desire a simple and convenient data product that can be expressed in Gaussian form as a mean

²For this example, we use WMAP7+BAO+ H_0 cosmology (Komatsu et al. 2011).

³Figure (5) can be directly compared to Figure 4 in Percival et al. (2010). However, we should point out that Percival et al. (2010) underestimate the tail of the likelihood, as they ignore the (small) possibility of misidentifying the BAO peak in their mock realizations.

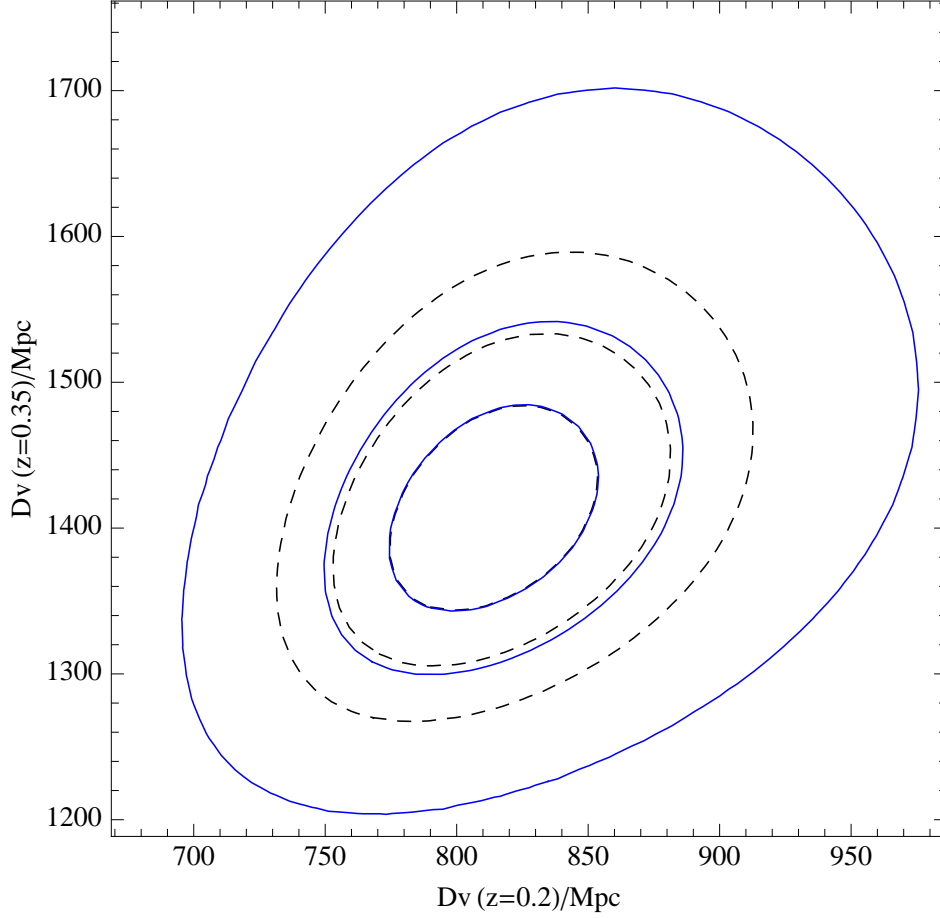


Fig. 5.— Contours of $\Delta\chi^2 = -2 \log[\text{likelihood}] = 2.3, 6.0$, and 11.8 (corresponding to 68.3%, 95.5%, and 99.7% gaussian probabilities) for distances to $z = 0.2$ and 0.35 , from BAO detection in SDSS7 galaxies (Percival et al. 2010). The dashed black contours are gaussian fits quoted in Percival et al. (2010), while solid blue contours correspond to our non-gaussian expectation, based on Eq. (3) (compare to Figure 4 in Percival et al. (2010)). Although the 68.3% (1σ) and 95.5% (2σ) confidence contours are essentially unchanged, the effect of the finite probability of having detected the BAO peak has a profound effect on the 99.7% (3σ) contour.

plus an error bar without using the full likelihood or redoing the initial analysis. This is exemplified by the many studies using the Gaussian approximation to the SDSS Baryon Acoustic Oscillation (BAO) distance.

In the case of marginal detections, such as in the case of BAO cosmology or gravitational wave astronomy in the near future, this will not suffice, and the possibility that the detections are pure noise, even if small, have to be taken into account to get both precise and accurate results. The main impact of taking this into account, is to radically alter the relationship between confidence intervals. While 68% confidence intervals may be essentially unchanged, 95%, 99.7% and higher confidence intervals can be radically altered. In the Gaussian approximation, these intervals are all trivially related to each other, but in the case where the finite detection probability is included this relationship is broken. Given that confidence limits must exceed 3σ to be generally accepted as providing truly secure knowledge about Nature, it is important to include finite detection probabilities in cosmological statistical analyses lest one’s posterior is stung by the tail of the scorpion. We provide an explicit and simple way to do this via Eq. (3), which uses the signal-to-noise of the detection to remove this sting.

A further concern might be that BAO analyses must assume a redshift-distance relation (typically Λ CDM) to compute the two-point correlation function in the first place, instead of recomputing $\xi(r)$ at every point in the cosmic parameter space. For models close to Λ CDM, this only causes small errors. However, if one wants to compute 3 - 5σ constraints one now includes models - because of the detection effect discussed in this *letter*- that are far from Λ CDM and where the approximation of using Λ CDM to convert from redshift to distance might be very poor. This raises additional issues about the true constraints on cosmic acceleration from current BAO surveys.

Finally, it is a well-known phenomenon (*e.g.*, Henrion & Fischhoff 1986) that experimentalists in all fields tend to underestimate the size of their error bars. This is particularly important here. Not only do the smaller error bars lead one to be over-confident about standard gaussian results, they are also likely to make one ignore the need for considering the finite-detection effects discussed here. For example, if one believes one has a 5 sigma detection of the BAO, then one could safely ignore the finite detection non-gaussianity. However, if one’s error bars were too small by a factor of two that would make the non-gaussian corrections much larger. In the case of BAO, estimating error bars on the correlation function is very difficult. Ideally one should run a very large number of mock simulations (*e.g.*, Takahashi et al. 2009), but in the absence of this it is common to use approximations such as jack-knife or log-normal simulations which may not give the correct answer.

The simplest way to take care of all of this effect is to use the full $P(k)$ power spectrum

measurements in Monte-Carlo Markov Chains, as is possible in *e.g.*, CosmoMC.

Acknowledgements: We thank Martin Kunz, Renée Hlozek, Hiranya Peiris and Will Percival for insights and particularly Chris Blake, who has always demanded 5σ for BAO detections. NA and BB are partially supported by the Perimeter Institute (PI). Research at PI is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research & Innovation. BB is supported by the South African National Research Foundation.

REFERENCES

- Bassett, B., & Hlozek, R. 2010, Baryon acoustic oscillations, ed. Ruiz-Lapuente, P., 246–+
- Blake, C., Davis, T., Poole, G. B., et al. 2011, MNRAS, 415, 2892
- Blake, C., Kazin, E. A., Beutler, F., et al. 2011, MNRAS, 1598
- Cai, Y.-C., Angulo, R. E., Baugh, C. M., Cole, S., Frenk, C. S., & Jenkins, A. 2009, MNRAS, 395, 1185
- Cole, S., et al. 2005, MNRAS, 362, 505
- Eisenstein, D. J., & Hu, W. 1999, ApJ, 511, 5
- Eisenstein, D. J., et al. 2005, ApJ, 633, 560
- Gaztañaga, E., Cabré, A., & Hui, L. 2009, MNRAS, 399, 1663
- Henrion, M., & Fischhoff, B. 1986, American Journal of Physics, 54, 791
- Hütsi, G. 2010, MNRAS, 401, 2477
- Kazin, E. A., Blanton, M. R., Scoccimarro, R., McBride, C. K., & Berlind, A. A. 2010a, ApJ, 719, 1032
- Kazin, E. A., et al. 2010b, ApJ, 710, 1444
- Kessler, R., Becker, A. C., Cinabro, D., et al. 2009, ApJS, 185, 32
- Komatsu, E., et al. 2011, ApJS, 192, 18
- Kunz, M., Bassett, B. A., & Hlozek, R. A. 2007, Phys. Rev. D, 75, 103508
- McDonald, P., & Eisenstein, D. J. 2007, Phys. Rev. D, 76, 063009

Miralda-Escude, J. 2009, ArXiv e-prints

Percival, W. J., Cole, S., Eisenstein, D. J., Nichol, R. C., Peacock, J. A., Pope, A. C., & Szalay, A. S. 2007, MNRAS, 381, 1053

Percival, W. J., et al. 2010, MNRAS, 401, 2148

Press, W. H. 1997, in Unsolved Problems in Astrophysics, ed. J. N. Bahcall & J. P. Ostriker, 49–60

Sawangwit, U., Shanks, T., Croom, S. M., et al. 2011, arXiv:1108.1198

Sullivan, M., Guy, J., Conley, A., et al. 2011, ApJ, 737, 102

Sylos Labini, F., Vasilyev, N. L., Baryshev, Y. V., & López-Corredoira, M. 2009, A&A, 505, 981

Takahashi, R., Yoshida, N., Takada, M., et al. 2009, ApJ, 700, 479

The Dark Energy Survey Collaboration. 2005, ArXiv Astrophysics e-prints

Zhan, H., Knox, L., & Tyson, J. A. 2009, ApJ, 690, 923